

Part II

Features

APEX Blind Deconvolution of Color Hubble Space Telescope Imagery and Other Astronomical Data

Over the last 50 years, signal recovery has become a fundamental and pervasive scientific activity. In research spanning length scales from the atomic to the cosmological, deconvolution now plays a critical role in extracting scientific content. In particular, because of its critical role in measurement science, the study of improved methods for deconvolution has been a long interest of MCSD. In recent years our attention has turned to the challenging problem of image deblurring, where we have developed a number of novel techniques. For example, NIST's APEX method is a technique for real-time blind deconvolution. It is blind in that the point-spread function (PSF) causing the blur is not known a priori. The PSF is taken from a class of 2-dimensional heavy-tailed probability density functions whose parameters are estimated from the image itself. Once the PSF is determined, the deconvolution proceeds by marching a diffusion equation backwards in time. The method relies on FFTs as its computational kernel, and hence is highly efficient, even for 1024×1024 images. The APEX method has been shown to be effective on a wide range of imagery, including medical images and scanning electron microscope images. Recently this method was extended to the sharpening of color imagery and to astronomical data.

Alfred Carasso

NIST's APEX blind deconvolution method, initially conceived from an abstract mathematical blueprint, has been found useful in sharpening a wide variety of images, including Hubble Space Telescope color imagery. This is discussed in a recent paper by A. S. Carasso in the October 2006 issue of *Optical Engineering*, the flagship research journal of the International Society of Optical Engineering (SPIE) [7]. The paper provides a detailed discussion of the application of the APEX method to astronomy, and includes visually striking enhancements of color Hubble imagery, as well as imagery from ground-based telescopes. One vivid example involving the Andromeda galaxy (Fig. 3) was selected as cover art by the journal Editors. Previous NIST work applying the APEX method to scanning electron microscopy (SEM) was published in the same journal some time ago [5]. Such microscopy work is ongoing, and remains of vital significance to NIST. Indeed, powerful new state-of-the-art SEM equipment has been acquired in support of NIST nanotechnology research. Considerable synergy is likely to be realized by comparing and contrasting the performance of the APEX method in such widely different contexts.



Figure 3. Kitt Peak true-color Andromeda galaxy image. Top: original image. Bottom: after APEX processing.

NASA's Hubble Telescope was launched in 1990, with an estimated total lifetime cost of around six billion dollars. Over 750,000 priceless images had been recorded and archived as of April 2005. These images have fundamentally increased our understanding of the universe, have enthralled the general public, and have significantly increased public awareness of scientific research.

Valuable additional information can quite likely be obtained from many of these images by the application

of appropriate sharpening methods. This is illustrated with the Tadpole galaxy example in Fig. 4. On the left is the original Hubble image taken in April 2002 using NASA's ultra sophisticated Advanced Camera for Surveys (ACS). Several clusters of young bright blue stars are visible in the Tadpole's spiral head and long tail. These and other foreground features are brought into sharper focus by applying the NIST-developed APEX sharpening procedure as shown on the right. However, of even greater interest, are the background features, which become much more evident in the APEX enhanced image. This background has been described by NASA as a "Whitman's Sampler" of galaxies, many of which are extremely far away, and represent "fossil samples of the universe's 13 billion year evolution."

Background. Deconvolution is the process of recovering a signal from measured data which has been subjected to a type of averaging process. For example, the fundamental 2-D deconvolution problem is to recover a signal f , such that

$$\int_{\mathbb{R}^2} h(x-u, y-v) f(u, v) du dv = g(x, y)$$

where g is the recorded data (which may contain noise) and h represents the averaging process. The study of deconvolution requires tools from several branches of mathematics, including integral equations, calculus of variations, operator theory, probability theory, Fourier analysis, partial differential equations, numerical analysis, linear algebra, and wavelet theory. Not surprisingly, MCS&D has long supported research in this area and is uniquely well-equipped for this task.

In prior years, our work focused on challenging 1-dimensional deconvolution problems of importance to the NIST measurement laboratories. These typically featured complex signals exhibiting non-differentiable singularities. Such behavior is associated with the arrival of propagating disturbances, and is of prime interest. One such situation occurs in acoustic emission studies, where experimental identification of elastic Green's functions is necessary. Another is the identification of dynamic Green's functions in complex structural networks, such as an orbiting space station. In each case, novel deconvolution techniques were devised by Carasso that could capture the singular Green's profile in the presence of noise [1, 2].

Much current scientific data is in the form of digital imagery. Whether the imaging instrument is an optical microscope, an electron microscope, a telescope, or any one of several medical imaging modalities, the resulting image is blurred due to distortions induced by the instrument, as well as by the medium through which the signal travels. In many cases, this distortion can be summarized and described by the imaging system's *point spread function* or PSF. Knowledge of the PSF, when this is available, enables sharpening of the

blurred image through very careful numerical treatment of the recorded image. That task is highly non-trivial, as deconvolution is an unstable computational process that can severely amplify noise. Fundamental work in this area was done by Carasso, who invented the *slow evolution from the continuation boundary* (SECB) constraint to stabilize the inversion process [3]. Subsequently, further significant work focused on finding the right mathematical function space in which to pose the deblurring problem. Indeed, the commonly used space of functions of bounded variation has the notorious property of eliminating fine scale details in an image, and such loss of texture is wholly unacceptable in practice. A cure for this so-called *staircase effect* was urgently needed. In 2003, Carasso showed that most natural images are not of bounded variation, but rather belong to Lipschitz spaces of non-differentiable functions. Carasso then used the Poisson singular integral to create an effective deblurring method that can recover texture [6]. U.S. patents have been obtained, or are in process, for these inventions.

Blind Deconvolution. In many cases, the PSF describing the blur is unknown or incompletely known. Blind deconvolution seeks to deblur the image without knowing the PSF. This is an extremely challenging problem in which severe ill-conditioning is compounded with non-uniqueness of solutions. Most known approaches are iterative in nature and seek to simultaneously reconstruct both the PSF and the deblurred image. Such iterations are typically ill-behaved and may develop stagnation points or diverge altogether. When the process is stable, many thousands of iterations and several hours of computing are often necessary to resolve fine detail. Such procedures are not feasible for real-time processing of large size images of complex objects.

Published in 2001, NIST's APEX method represents a major breakthrough in blind deconvolution [4]. Like his previous work, Carasso's research in blind image deblurring has focused on developing reliable direct (non-iterative) methods, in which fast Fourier transform (FFT) algorithms are used to solve appropriately regularized versions of the ill-posed deblurring problem. The SECB and Poisson singular integral methods are direct FFT methods that assume the PSF is known. The blind APEX method is predicated on a preselected class of blurs in the form of 2-D radially symmetric, bell-shaped, heavy-tailed probability density functions, a class that generalizes Gaussian and Lorentzian distributions. By design, these blurs are chosen so as to have a crucial additional mathematical property, namely they must form a semigroup under convolution. An explicit mathematical formula for such blurs was worked out by the legendary French probabilist Paul Lévy in the 1930's in his celebrated work of the Central Limit Theorem. Use of that formula enables initial trial identification of the blur parameters by Fourier analysis

of the blurred image. Naturally, there is no assurance that this class of blurs is appropriate for a given image, and, in fact, not all images can be usefully enhanced with the APEX method. Nevertheless, APEX is effective on a surprisingly wide variety of images.

The semigroup property allows reformulation of the deconvolution problem into a mathematically equivalent problem of solving a generalized diffusion equation backwards in time, given the blurred noisy image as data at time $t=1$. Here, SECB regularization enables stable marching backwards towards the desired deblurred image at time $t=0$. Such *slow motion* deconvolution is a vital element of the APEX method. It allows for visual monitoring of the process, for the calculation of statistical diagnostics, and for the opportunity to curtail the process when noise amplification and/or ringing artifacts begin to appear. Such early termination is equivalent to interactive readjustment of the initially detected blur parameters. Because the entire computation is FFT-based, high resolution images, of size 1024x1024 pixels, can be handled in about 2-3 minutes on current desktop computers. Numerous trial deconvolutions using different parameters can thus be done. Such efficient exploration of parameter space is often a key to the successful solution of ill-posed inverse problems.

Astronomical color imagery. Blind deconvolution of astronomical images is a challenging problem of considerable interest. The applicability of the APEX method is not evident a-priori, and hinges on whether the assumed class of heavy-tailed blurs is appropriate for the various types of modern astronomical imaging devices. Color imagery presents a significant additional challenge, as each color component may have its own individual PSF, which must be identified separately. As is clear from Figs. 3 and 4, color is an important element in astronomical imagery. There is the real danger that the deconvolution process may upset the original balance of colors, and result in a deblurred image with physically false colors. Blind deblurring of color imagery is a subject that is still very much in its infancy.

A remarkable property of the APEX method is the ease with which it can be applied to color imagery, and the plausibility of the ensuing results [7]. This property stems directly from its ability to permit monitoring of the unfolding process in slow motion. In addition,

quantities such as the image L_1 norm and the image gradient norm can be calculated at each discrete time level on the way to $t=0$. In well-behaved deconvolution, the L_1 norm of each color component should be conserved, while the gradient norm should increase monotonically. One can enforce L_1 norm conservation in each color component separately by terminating the deconvolution process for that component whenever the current L_1 norm exceeds its original value at $t=1$ by more than a very few percent. This strategy of enforcing componentwise L_1 norm conservation has been found to maintain the balance of colors in all of the color examples to which the APEX method has been applied.

The example in Fig. 3 is one where the APEX-detected blue component PSF was much narrower than the red and green PSFs, indicating that the blue component image was perceived by APEX to be less blurred than the red and green. Such differential deblurring has clearly preserved the balance of colors.

References

- [1] A. Carasso and N. N. Hsu, Probe Waveforms and Deconvolution in the Experimental Determination of Elastic Green's Functions, *SIAM Journal on Applied Mathematics* **45**, (1985), pp. 369-383.
- [2] A. Carasso and E. Simiu, Identification of Dynamic Green's Functions in Structural Networks, *AIAA Journal* **27**, (1989), pp. 492-499.
- [3] A. Carasso, Overcoming Hölder Continuity in Ill-posed Continuation Problems, *SIAM Journal on Numerical Analysis* **31** (1994), pp. 1535-1557.
- [4] A. Carasso, Direct Blind Deconvolution, *SIAM Journal on Applied Mathematics* **61** (2001), pp. 1980-2007.
- [5] A. Carasso, D. Bright, and A. Vldar, APEX Method and Real-time Blind Deconvolution of Scanning Electron Microscope Imagery, *Optical Engineering* **41** (2002), pp. 2499-2514.
- [6] A. Carasso, Singular Integrals, Image Smoothness, and the Recovery of Texture in Image Deblurring, *SIAM Journal on Applied Mathematics* **64** (2004), pp. 1749-1774.
- [7] A. Carasso, APEX Blind Deconvolution of Color Hubble Space Telescope Imagery and Other Astronomical Data, *Optical Engineering* **45**, October 2006, 107004, 15 pages.

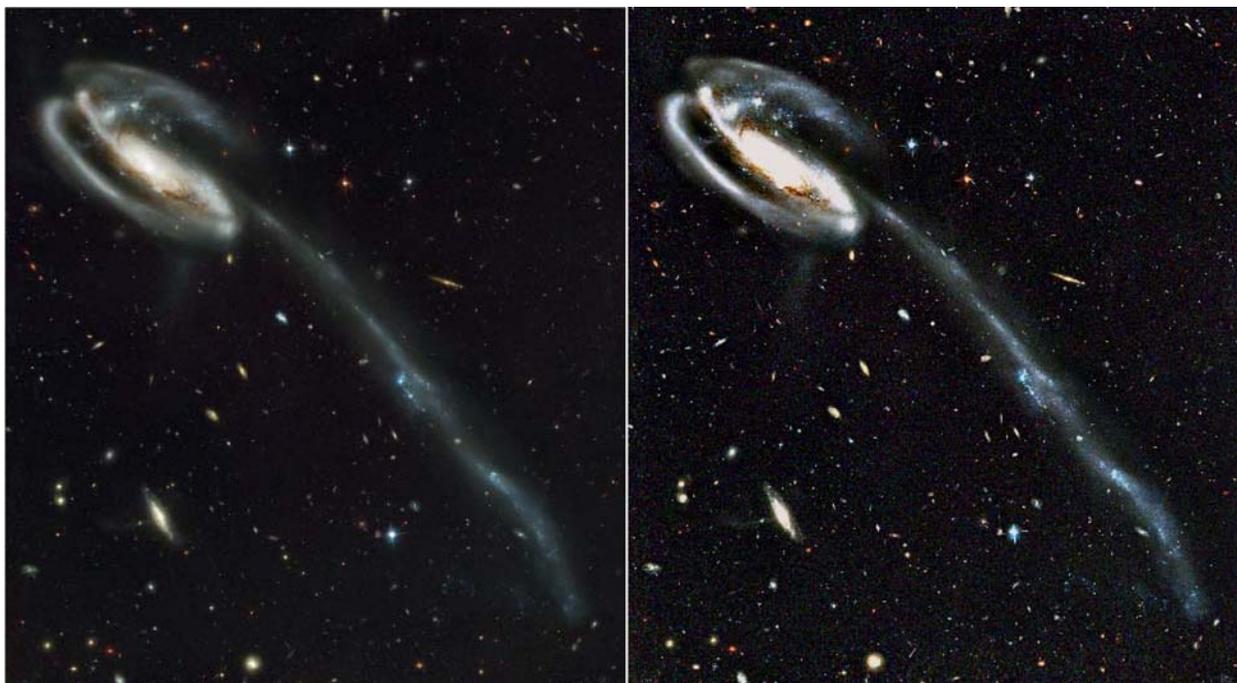


Figure 4. APEX processing of Hubble TADPOLE galaxy. Enhancement of background is of considerable interest. Background said to contain a "Whitman sampler of galaxies stretching back to the beginning of time. Image obtained by NASA's most sophisticated Advanced Camera for Surveys (ACS).

Optical Coherence Tomography for Biomedical Imaging

A wealth of biomedical imaging techniques are being developed for diagnosis and monitoring of a wide range of diseases. Well-known modalities include MRI, PET scans, and quantitative x-ray tomographic technologies (qCT). Recently, interest in optical coherence tomography (OCT) has grown as it has been demonstrated to give quantitative structural information pertaining to biological scatterers. One proposed application has been to use OCT as a diagnostic tool for cancer. NIST researchers are performing quantitative evaluations of this imaging modality. In doing so, a new channel for potentially useful diagnostic information has been identified.

Andrew Dienstfrey

Over a range of optical frequencies and scattering parameters, it has been shown that cell nuclei are the dominant contributors to back-scattered radiation. Thus, measurements can be inverted for nuclei features, for example the characteristic size. As many epithelial cell cancers (e.g., cervical, esophageal, colon, skin, oral) exhibit enlarged cell nuclei as a precancerous condition, it is hoped that a fiber-based OCT technique could serve as an early diagnostic tool for these disease models. Additional biomedical uses of OCT currently being investigated include: cellular substructure imaging, assessment of the efficacy of chemopreventative agents, and the measurement of characteristics of large-scale intracellular organization.

One of the novel characteristics of OCT is that both the intensity and phase of the fields of interest can be measured. This ability, although not unique to OCT, nevertheless distinguishes it from the majority of optical techniques which measure only the intensity (magnitude) of the optical field. Four OCT platforms were built by Shellee Dyer, Tasshi Dennis, and Paul Williams in NIST's Opto-electronics Division to investigate the quantitative potential of this technology.

For fixed polarization and point in space, the relationship between the incident and scattered fields may be viewed as a linear dynamical system. For such systems we have $E_s(f) = E_i(f) H(f)$ where f is the frequency of incident radiation, $E_{s,i}(f)$ are polarizations of scattered and incident fields, and $H(f)$ is the Fourier transform of the impulse response function. In some cases, measurements of $H(f)$ can be inverted for physical characteristics of the scatterers. As $H(f)$ is complex, its complete characterization requires the optical phases of $E_{s,i}(f)$. In addition to the complication of measuring optical phase, some investigators have assumed that this quan-

tity carries no information not already encoded by the magnitude. Writing $H(f) = \rho(f) \exp(i\phi(f))$, the mathematical statement is that $H(f)$ is minimum phase, in which case $\phi = K \ln(\rho)$, where K is a well-known singular integral operator. This minimum-phase assumption is common but, to our knowledge, has not been sufficiently explored until now.

A classical structure theorem from linear systems theory is a powerful aid in the analysis. The derivative of the phase with respect to frequency, $\phi'(f)$ is known as the relative group-delay (RGD). In terms of the RGD the theorem states that

$$\Delta'(f) \equiv \phi'(f) - \phi'_{mp}(f) = 2\pi\tau + \sum_n^N \frac{\gamma_n}{(f - f_n)^2 + \gamma_n^2}$$

In this equation τ denotes a time shift, and the summation terms come from "all-pass" or "Blaschke factors." Details aside, the conclusion is that the discrepancy between a minimum and non-minimum-phase response is not arbitrary but, rather, has a structure dictated by general theory.

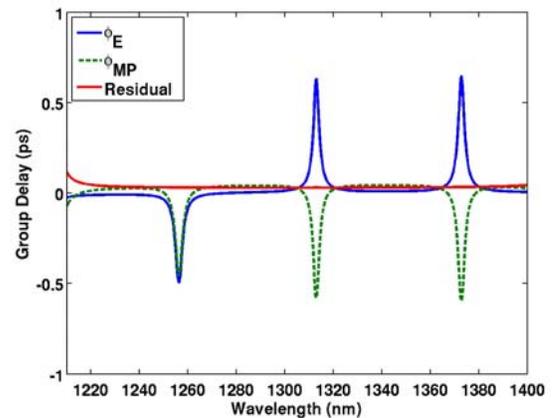


Figure 5. True group delay and minimum phase group delay (blue solid and green dashed curves) are plotted as a function of wavelength for a single sphere ($r=5$ microns, $n=1.41+0.003i$). The true group delay resonances at $\lambda=1.31\mu$ and $\lambda=1.375\mu$ are inverted from those predicted by minimum phase analysis. Correcting the minimum phase for time shift and Blaschke factors as predicted by theory brings the two RGD in agreement. The residual after correction is shown in red.

Two suites of software tools applicable to this problem have been developed within MCS D. The first computes full wave electromagnetic solutions to Maxwell's equations with multiple, fully general spherical scatterers in arbitrary position. The algorithm represents electromagnetic fields in terms of vector wave functions centered at sphere locations. Coupling between mutual spheres is achieved using appropriate translation operators. The second suite of tools performs the minimum

phase analysis. These codes compute the singular Kramers-Kronig operator K via product quadrature of a spline representation of sampled magnitude data. All algorithms were verified for accuracy and convergence by extensive numerical testing. Additionally, the spherical scattering code was validated by comparison with a single sphere code (Mie theory) developed previously by investigators in the NIST Physics Lab.

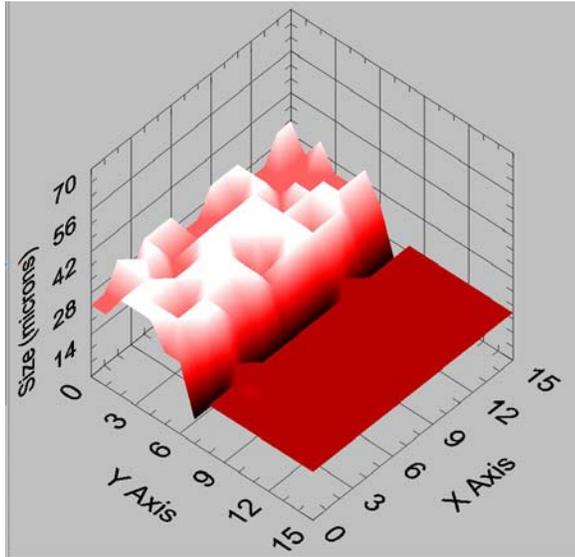


Figure 6. Characteristic size of polystyrene spheres adhered to glass coverslip as determined by inversion of OCT measurements of back-scattered group delay. The ridge indicates the boundary between small and large spheres ($r=14$ and 28 micrometers). Spatial averages as a function position are indicated by the z-axis. The sphere positions and radii correspond well with measurement by microscopy.

These codes were used to compute vector fields scattered by isolated spheres, along with the true complex response function, $H(f)$, as a function of frequency. The minimum phase response $\phi_{mp}(f) \equiv [K \ln(|H|)](f)$ was compared with the true phase response $\phi_E(f) \equiv \text{Arg}(H(f))$. The group delays (derivatives) are plotted in Fig. 5. At the resonance near $\lambda=1250\text{nm}$, the true and minimum RGD agree. The two other resonances, however, are of opposite signs. According to theory, the RGD difference must be represented by a sum of a time-shift and Blaschke terms. A non-linear fit for τ and parameters $\{f_n, \gamma_n\}$ was performed. The minimum phase RGD augmented by the exponential and Blaschke terms determined by the non-linear analysis is subtracted from the true phase, and the difference (residual) is also shown in the figure. Clearly the differences between the minimum and true RGD are accounted for by the Blaschke terms as predicted by theory. This result demonstrates that back-

scattering by a single dielectric sphere can exhibit both minimum and non-minimum-phase response. Although of an elementary character, to our knowledge this observation has not been made before. The conclusion is surprising as it contradicts heuristic notions of minimum-phase held by many in the optics and bioimaging communities.

Supported by these results, we find that optical phase offers potentially new diagnostic information independent of intensity. Additionally, for quantities that admit derivation by independent analysis of either phase or intensity, the error characteristics need not be the same. For example, many algorithms exist to invert intensity for scatterer size. Such procedures are prone to a systematic error in that the intensity of broadband light sources generally is not constant as functions of frequency. This modulation must be deconvolved from the experimental measurement so as to isolate the effect due to scattering. By contrast, phase analysis is invariant with respect to the intensity provided only that the intensity is non-zero. Thus phase-based analyses are self-correcting with respect to this form of error. A microscope slide was coated with polystyrene spheres drawn from two populations. The manufacturer specified the average radii as $r_1=15\mu$ and $r_2=26\mu$, both with 15% variance. The slide was imaged using the NIST spectroscopic OCT platform, and group delay measurements inverted for characteristic sphere size as a function of xy -location on the slide. The results are shown in Fig. 6. Quantitative comparisons of the radii computed by OCT compare well with independent measurements made via optical microscopy, as well as with the manufacturer's specifications.

References

- [1] T. Dennis, S. D. Dyer, and A. Dienstfrey, Phase Dispersion Light-scattering for Quantitative Size-imaging of Spherical Scatterers, *Optics Letters*, submitted.
- [2] S. D. Dyer, T. Dennis, L. K. Street, S. M. Etzel, T. A. Germer, and A. Dienstfrey, Spectroscopic Phase-Dispersion Optical Coherence Tomography Measurements of Phantoms, *Optics Express* **14** (18), 8138–8153, Aug. 2006.

Participants

A. Dienstfrey (MCSD); S. Dyer, T. Dennis, and P. Williams (EEEL)

Measurement and Analysis of Tissue Engineering Scaffolds

Advances in tissue engineering are leading to the ability to grow cells for use in repairing or replacing skin, bones, and parts of bodily organs. In such applications, cells are seeded onto artificial structures, called scaffolds, where they undergo growth. Success of such procedures requires a quantitative understanding of the underlying growth media and its impact on cell growth. To this end, scientists in NIST's Materials Science and Engineering Laboratory (MSEL) have gathered data on tissue scaffold materials using a variety of imaging techniques. We are working with MSEL scientists to bring multi-modal imaging data on tissue engineering scaffolds into a virtual environment, where interactive measurement and analyses may be done. By developing techniques and tools to support applications of this type, we are extending the use of visualization from the qualitative to the quantitative.

John Hagedorn

One important consideration for the viability of tissue engineered products is the relationship between cell proliferation and the microstructure of the underlying scaffold material. Scientists in the NIST Materials Science and Engineering Laboratory have developed techniques for generating three-dimensional images of such scaffolds using X-ray micro-computed tomography (μ CT). Such methods have diverse application in areas such as the failure analysis of polymer composites and the reliability of semiconductor devices. A recent application has been to the development and characterization of standard reference materials (SRMs) for the growth of tissue engineered products. SRMs of this type are needed for industry to develop low cost manufacturing processes required for commercial success.

We are pursuing the use of the immersive visualization environment (IVE) as a framework for more easily measuring distinguishing properties of scaffolds, for support in developing consensus definitions of scaffold descriptors, for understanding a variety of descriptor measurement methods, and for qualitatively evaluating and validating scaffold manufacturing techniques. The use of such real time immersive visualization (IV) techniques has been enabled by the continuing acceleration in speed and capability of commodity graphics processors. For this application, we have developed software that enables the measurement and analysis of tissue engineering scaffold materials from three-dimensional μ CT data, which is segmented and converted to a polygonal representation for use in the IVE. In addition, we have developed software enabling a

range of interactive measurements to be performed on such data in the virtual environment.



Figure 7. An IVE user making scaffold measurements.

For example, this software enables us to compare an idealized scaffold with a manufactured scaffold to determine differences in scaffold strut properties. Essential to this work is the use of immersive visualization which gives the researcher the ability to interact directly with data representations in ways that are not possible during data acquisition or with desktop systems. Structures are inspected and measurements made and analyzed during the immersive session. Using these measurements, researchers can assess the fidelity of actual scaffolds to the design model and evaluate the scaffold manufacturing processes.

While the initial processing of the image data was relatively straightforward, the analysis and measurement of geometric descriptors was more challenging. We built a software system within the IV environment that integrates the interactive measurement of scaffold characteristics, the analysis of the collected measurements, the display of the analysis, and the interactions with the data and analyses that enable grouping of results. To accomplish this, we implemented three types of measurement capabilities: measurement with lines, with cylinders, and with ellipsoids. The software enables the user to shape and place these objects interactively in the immersive environment. These forms are matched to the forms of fibers and pores in the data and serve as measuring devices. Analyses of the created geometry are presented to the user, and the user can interact with the data to derive additional information.

The initial measurement task that we undertook was the manual measurement of linear distances. It was felt that this step would enable the understanding of several important scaffold characteristics, one of which is strut diameter distribution and any associated anisotropy.

The specific scenario for this first implementation would be: (1) The user interactively collects a set of linear measurements. (2) A statistical analysis is made. (3) The analysis, including the distribution of measurements, is presented to the user. (4) The user interacts with the measurement distribution in order to highlight measurements that fall within any selected range of values. All of these tasks would be performed in real-time during the IV session.

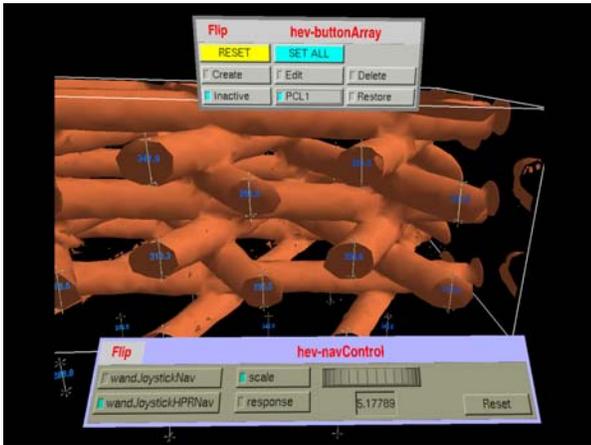


Figure 8. Scaffold cross sections with diameter measurements.

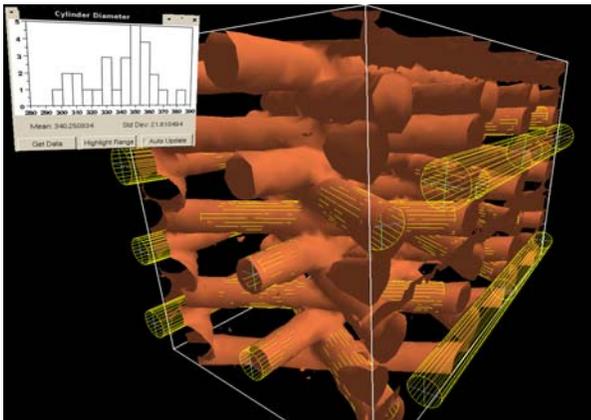


Figure 9. Cylinder measurement and analysis tool.

We performed measurements both on the “as-designed” scaffold model (generated synthetically from the design) and on the image of the actual manufactured scaffold material. The former are intended to validate the measurement method, while the latter are used to understand the scaffold structural characteristics and fabrication method. We derived descriptors such as gap width, spacing of fibers, angles between intersecting struts, etc., and compared the results to those from the as-designed scaffold. We found that the inter-junction strut diameter was about 19% smaller than the as-designed model. The at-junction strut diameter (or layer thickness) was about 33% smaller than the as-designed model. But the angular measurements correspond very closely to the design.

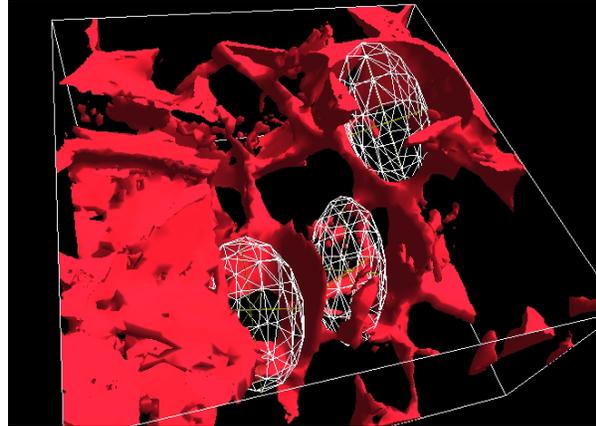


Figure 10. Ellipsoid measurement tool used on a scaffold.

We have found that IV enables both qualitative and quantitative understanding of 3D structure of tissue engineering scaffolds not otherwise possible. Measurements made within the virtual environment would have been very difficult to make with typical desktop visualization techniques.

The following IV measurement of descriptors should be implemented to enable further evaluation: layer planarity, strut diameter uniformity, strut circularity, and strut location. We also plan to use the immersive environment in conjunction with automatic measurement techniques that we are developing to aid in understanding the action of the automatic algorithms and as a way of validating those methods.

References

- [1] J. Hagedorn, J. Dunkers, A. Peskin, J. Kelso, J. Terrill, Quantitative, Interactive Measurement of Tissue Engineering Scaffold Structure in an Immersive Visualization Environment, *Biomaterials Forum* **28**:4 (2006), pp. 6-9.
- [2] J. Dunkers, J. Hagedorn, A. Peskin, J. Kelso, J. Terrill, and L. Henderson, “Interactive, Quantitative Analysis Of Scaffold Structure Using Immersive Visualization,” in *Proceedings of the 2006 Summer Bioengineering Conference*, Amelia Island Plantation, Amelia Island, FL, June 21- 25, 2006.
- [3] J. G. Hagedorn, S. G. Satterfield, J. T. Kelso, W. Austin, J. E. Terrill, A. Peskin, Correction of Location and Orientation Errors in Electromagnetic Motion Tracking, *Presence*, to appear.

Participants

J. Hagedorn, J. Kelso, A. Peskin, S. Satterfield, J. Terrill (MCS D), J. Dunkers, L. Henderson, M. Cicerone, L. Levine (NIST MSEL).

<http://math.nist.gov/mcsd/savg/vis/tissue>

Visualization of Complex Function Data

Effective visualizations can help researchers obtain a more complete understanding of high level mathematical functions that arise in many applications. Thus, dynamic interactive 3D graphs of function surfaces will be a key component of the NIST Digital Library of Mathematical Functions, which is being developed to replace the classic NBS Handbook of Mathematical Functions (Abramowitz and Stegun, 1964). Designing software to plot complicated 3D surfaces can be a challenging task. The function data must be computed accurately, the plot must capture important surface features with little user input, and the plot should be accessible to users on a variety of platforms. Using techniques from numerical grid generation and various technologies, including the Virtual Reality Modeling Language, X3D, embedded 3D, and video capture, we have been able to address most of these issues to produce precise and informative visualizations.

Bonita V. Saunders

The NIST Digital Library of Mathematical Functions (DLMF) will include formulas, methods of computation, graphs, references, and links to software for more than 40 classes of high level mathematical functions. The web site will also feature interactive navigation, equation search, and dynamic interactive visualizations. A sample chapter on the gamma function is available at <http://dlmf.nist.gov/Contents>.

Precise 3D visualizations of complex functions enable exploration of attributes such as poles, zeros, and branch cuts. The development of such visualizations requires the accurate computation of function data and the determination of the best methods for rendering and disseminating the visualizations in an environment accessible to users on various platforms.

Function Evaluation. A key concern when constructing plots of complicated mathematical functions is accuracy, in both the data accuracy and the plot itself. To validate data accuracy we compute function values using at least two independent methods. This might involve the use of standard computer algebra packages, routines from commercial libraries and free repositories, or even personal Fortran and C codes from the chapter authors. While authors may provide initial versions of the plots, all functions are recomputed at NIST to insure the validity of the data.

Plot accuracy concerns the visual display of the data. Accurate visual representation of the function depends not only on the accuracy of the data, but also on the plotting tool or package. Commercial packages often

have many built-in special functions, but their 3D plots are usually over a rectangular mesh, often leading to poor or misleading graphs. When function values lie outside the range of interest, many packages have trouble properly clipping the function surface. Furthermore, even when the plot looks satisfactory inside a package, it may be completely unacceptable when transformed for display on the web.

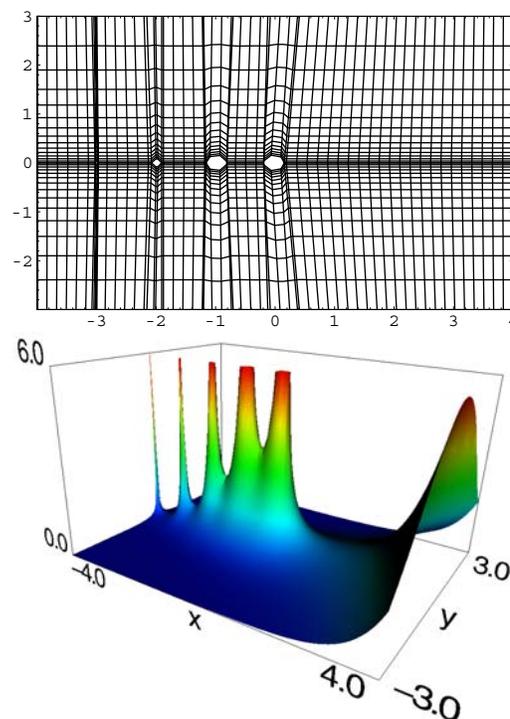


Figure 11. Computational grid and modulus of the complex Gamma function $\Gamma(x+iy)$.

We have solved many of these problems using techniques of numerical grid generation, such as transfinite blending function interpolation and a modified tensor product spline generator, to design customized meshes fitted to selected contours of the function. By computing the function values over such a mesh, we can accurately represent key function features. See Fig. 11.

VRML/X3D. Once the function data is computed, it must be translated into a format that allows viewing on the web. We use Virtual Reality Modeling Language (VRML), a standard 3D format for which browsers and plugins for a variety of platforms are publicly available. While standard VRML controls such as rotate, zoom, and pan permit a user to examine a 3D display from an arbitrary direction and position, we have implemented additional capabilities such as dynamic cutting plane control, color map control, and scale control. For example, cutting plane control allows one to exam-

ine the intersection of a plane with a function surface as the plane moves through the surface in a direction parallel to the x , y , or z coordinate axis.

Many DLMF visualizations represent real or complex valued functions of the form $z=f(x,y)$. Users may choose a height-based color map, where height is $|z|$ if f is complex. For complex-valued functions one may also choose a mapping based on the phase of z . Two options for the phase based mapping are available: a continuous mapping of phase to color or a four-color mapping based on the quadrant of the phase angle. Fig. 12 illustrates the continuous spectrum mapping. Fig. 13 illustrates use of the scale control.

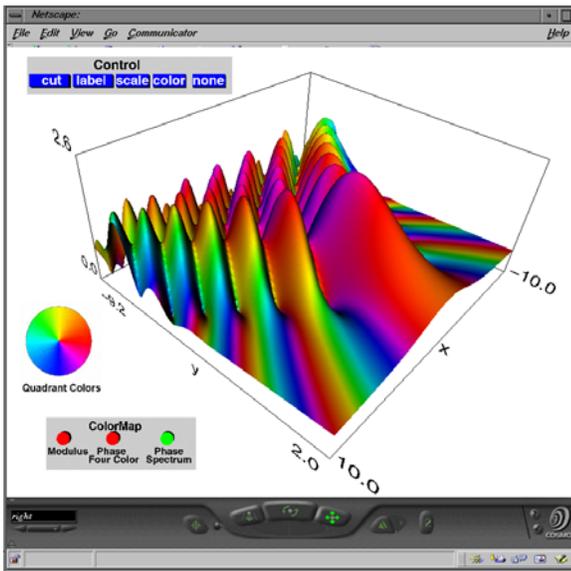


Figure 12. Modulus of complex Pearcey integral function $\Psi_2(x,y)$ with continuous spectrum color map.

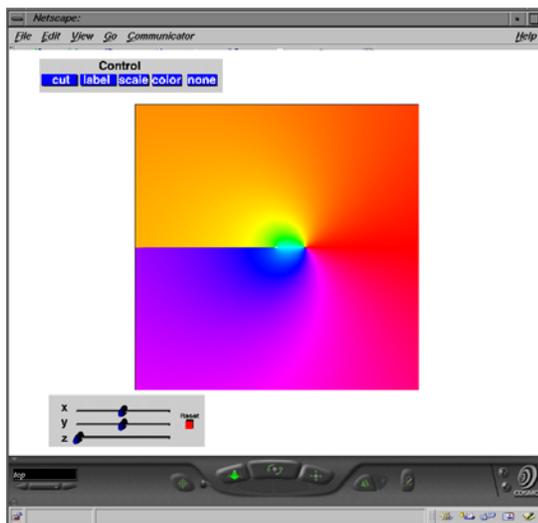


Figure 13. Surface scaled down in z direction. Modulus of the complex log function is shown with phase based color map. Note the branch cut along the negative x axis.

Alternative Technologies. In spite of the availability of VRML browsers, several things might preclude some users from gaining access to the DLMF visualizations. For example, a browser might not be available for their particular platform or configuration, or security concerns at their site may prevent the download of new software. To address such concerns, we have investigated several alternative technologies to render and disseminate the visualization of complex mathematical functions on the web: PDF document with embedded 3D graphics, animated movies using Apple QuickTime VR, and animated movies using video capture software. Based on user interaction capabilities and platform portability, it appears that embedded 3D graphics technology offers the most promise. We have used Adobe Acrobat 3D to create prototype interactive visualizations in PDF.

Over one hundred Web-based 3D visualizations for the NIST DLMF have been completed using VRML. Feedback from DLMF editors, authors, and other observers has been extremely positive. We continue to improve existing features and add new ones. One new feature being tested, for example, is a user capability to click anywhere on the function surface and obtain the coordinates. We also are seeking ways to improve the performance of the cutting planes, and working on adaptive grid techniques to decrease the size of some of our computational meshes.

References

- [1] Q. Wang and B. Saunders, Web-Based 3D Visualization in a Digital Library of Mathematical Functions, in *Proceedings of Web3D 2005: 10th International Conference on 3D Web Technology*, Bangor, Wales, United Kingdom, 2005.
- [2] B. Saunders and Q. Wang, Boundary/Contour Fitted Grid Generation for Effective Visualizations in a Digital Library of Mathematical Functions in *Proceedings of the Ninth International Conference on Numerical Grid Generation in Computational Field Simulations*, San Jose, California, 2005.
- [3] Q. Wang, B. Saunders and S. Ressler, Dissemination of 3D Visualizations of Complex Function Data for the NIST Digital Library of Mathematical Functions in *Proceedings of 20th International CODATA Conference, Beijing, China : Special Issue of CODATA Data Science Journal*, 2006, to appear.

Participants

B.V. Saunders (MCSD), Q. Wang, S. Ressler (ITL), D. Lozier, F. Olver (MCSD).

High Precision Calculation of Fundamental Properties of Few-Electron Atomic and Molecular Systems

Computational scientists at MCSD and Indiana University have achieved record levels of accuracy in the development of computational methods for determining fundamental properties of molecules. Their recent result for the ground state of dihydrogen (H_2) represents the highest level of accuracy ever reached (10^{12} hartree) in molecular quantum computations (except for trivial one-electron cases). H_2 has the distinction of being the first molecule whose dissociation energy was correctly predicted by quantum mechanical calculation (1968) before being measured reliably by experiment. Today we may be witnessing again a situation in which quantum mechanical calculations (“virtual measurement”) yield more accurate determinations of this fundamental property than can be measured experimentally.

James Sims

NIST has long been involved in supplying critically-evaluated data on atomic and molecular properties such as the atomic properties of the elements contained in the Periodic Table and the vibrational and electronic energy level data for neutral and ionic molecules contained in the NIST Chemistry WebBook. Fundamental to this endeavor is the ability to predict, theoretically, a property more accurately than even the most accurate experiments. It is our goal to be able to accomplish this for few-electron atomic and molecular systems.

Impressive advances have been made in the study of atomic and molecular structure in both theory and experiment. For atomic hydrogen and other equivalent two-body systems, exact analytical solutions to the nonrelativistic Schrödinger equation are known. It is now possible to calculate essentially exact nonrelativistic energies for atomic helium (He) and other three-body systems as well. Even for properties other than the nonrelativistic energy, the precision of the calculation has been referred to as “essentially exact for all practical purposes” [1], i.e., the precision goes well beyond what can be achieved experimentally. Notwithstanding this theoretical advance, the scarcity of information (experimental and theoretical) on atomic energy levels is overwhelming, especially for highly ionized atoms. On the theoretical side, the availability of high precision results tails off as the state of ionization increases, not to mention higher angular momentum states. In addition, atomic anions have more diffuse electronic distributions, and therefore represent more challenging computational targets. On the dia-

atomic molecular side, the presence of two fixed nuclei rather than one makes the integrals harder, hence exact analytical solutions to the molecular nonrelativistic Schrödinger equation are known (in the Born-Oppenheimer approximation) only for the one-electron H_2^+ ion and other equivalent one electron systems. The challenge for computational scientists is to extend the phenomenal accomplishments on atomic helium to three, four, and more electron atomic states and to molecular systems.

Dihydrogen, H_2 , has the distinction of being the first molecule whose dissociation energy was correctly predicted by quantum mechanical calculation (1968) before this same quantity was measured reliably by experiment. Today we may be witnessing again a situation in which quantum mechanical calculations yield more accurate determinations of this fundamental property than can be measured experimentally. Very high precision approximations are now available for molecular hydrogen (a two-electron system) as a result of our calculated Born-Oppenheimer energies of $^1\Sigma_g^+$ states of H_2 using up to 7034 expansion terms in confocal elliptical coordinates with explicit inclusion of interelectronic distance coordinates up through r_{12}^7 . We calculated Born-Oppenheimer (BO) energies for various internuclear distances in the range of 0.4 bohr to 6.0 bohr with an error of 1 in the 13th digit, for example, the nonrelativistic energy is -1.1744 7593 1399(1) hartree at the equilibrium $R = 1.4011$ bohr distance. The BO energies are what are used to provide the $X^1\Sigma_g^+$ potential energy curve which is critical to determining, theoretically, the dissociation energy D_0 of H_2 , as can be seen from the figure.

Our results [2] are the most accurate energy values ever obtained for a molecule of that size; 100 times better than the best previous calculated value or the best experimental value, but the methods used are perhaps equally important. The calculation requires solving an approximation of the Schrödinger equation, one of the central equations of quantum mechanics. It can be approximated as the sum of an infinite number of terms, each additional term contributing a bit more to the accuracy of the result. For all but the simplest systems or a relative handful of terms, however, the calculation rapidly becomes impossibly complex. While very precise calculations have been done for systems of just three components such as helium (a nucleus and two electrons), we are the first to reach this level of precision for H_2 with two nuclei and two electrons.

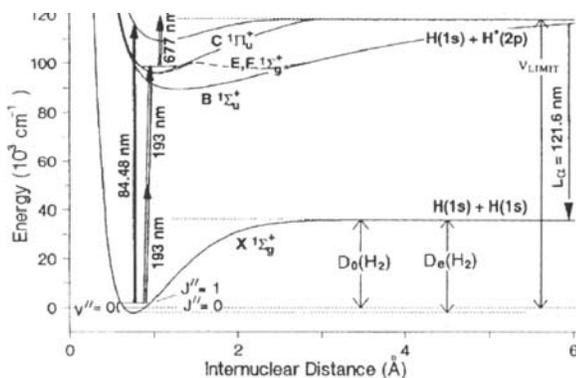


Figure 14. Potential energy curves of H_2 showing transitions from the ground state to the $B^u \Sigma_u$ state involved in photodissociation of H_2 .

To make the problem computationally practical, we merged two earlier algorithms for these calculations — one which has advantages in ease of calculation, and another which more rapidly achieves accurate results — into a hybrid with some of the advantages of both. We also developed improved computer code for a key computational bottleneck (high-precision solution of the large-scale generalized matrix eigenvalue problem) using parallel processing. Our final calculations were run on a 147-processor parallel cluster at NIST over the course of a weekend — on a single processor it would have taken close to six months.

Almost all results reported in our paper were obtained using quadruple precision (~30+ digits) floating point subroutines written in Fortran 90. In addition to quadruple precision, multiple precision floating point arithmetic was used for the R udenberg ϕ function, on which all integrals depend, and which is subject to catastrophic cancellations when computing differences in low precision. To address these problems we systematically increased the number of decimal digits used for only the ϕ part of the calculation up to a maximum of 160 decimal digits. In each case we obtained lower (i.e., more accurate) energies than previously reported results. The result for the ground state of H_2 represents, in fact, the highest level of accuracy ever reached (10^{-12} hartree) in molecular quantum computations (except for trivial one-electron cases) and is good enough to pin down the dissociation energy D_0 for H_2 to 0.001 wave number *if* one can separately calculate the adiabatic, relativistic, and radiative (QED) corrections accurately enough. This level of theoretical accuracy suggests that perhaps the experiment should be redone using 0.001 wave number as the target level of accuracy. That is, two orders of magnitude reduction in the current experimental uncertainties.

Now that high precision calculations are feasible for the two-electron molecular problem, we are shifting our focus back to the atomic problem. In going from

He (two electrons) to Li (lithium, three electrons) to Be (beryllium, four electrons), the situation vis a vis high precision calculations degrades to the point that already at four electrons (Be) there are no calculations of the ground or excited states with an error of less than 10^{-6} a.u. To alleviate this problem, we are currently building the infrastructure for a calculation on the ground state energy level of the Be atom. This has involved removing several bottlenecks to highly accurate calculations using our method, including treatment of the three electron triangle integrals [3], treating the remaining three-electron kinetic energy and nuclear attraction integrals [4], and dealing with four-electron integrals [5]. The beryllium calculation will determine whether highly accurate calculations on atoms with $N \geq 4$ can become a reality.

References

- [1] G. W. F. Drake and W. C. Martin, Ionization Energies and Quantum Electrodynamical Effects in the Lower $1s$ and $1snp$ Levels of Neutral Helium ($^4\text{He I}$), *Canadian Journal of Physics* **76**, 679–698 (1998).
- [2] J. S. Sims and S. A. Hagstrom, High Precision Variational Calculations for the Born-Oppenheimer Energies of the Ground State of the Hydrogen Molecule, *Journal of Chemical Physics* **124**, 094101 (2006).
- [3] J. S. Sims and S. A. Hagstrom, Mathematical and Computational Science Issues in High Precision Hy-CI Variational Calculations I. Three-electron integrals, *Journal of Physics B: Atomic Molecular and Optical Physics* **37**, 1519–1540 (2004).
- [4] J. S. Sims and S. A. Hagstrom, Mathematical and Computational Science Issues in High Precision Hy-CI Variational Calculations II. Kinetic Energy and Nuclear Attraction integrals, (2006), to be submitted.
- [5] J. S. Sims and S. A. Hagstrom, Mathematical and Computational Science Issues in High Precision Hy-CI Variational Calculations III. Four-electron Integrals, (2007), to be submitted.

Participants

J. Sims (MCSD); S. Hagstrom (Indiana University)

<http://math.nist.gov/mcsd/savg/parallel/atomic>

Benchmarks for Quantum Computing with Ion Traps

Quantum computers have the potential for significantly speeding up many useful algorithms. However, building quantum computers is challenging. Currently available quantum devices can realize computations with only a few quantum bits (qubits) and steps. Atomic qubits in ion traps are currently considered one of the leading candidates for realizing large quantum computers. To determine the potential of current ion trap quantum computers, we implemented two benchmark quantum subroutines for up to six atomic qubits and analyzed the error of the implementations. The first benchmark involved creating “Schrödinger cat states” of four to six qubits. The second demonstrated an entanglement purification protocol that is expected to play a key role in large scale implementations of quantum computers and quantum communication protocols.

Emanuel H. Knill

The goal for experimental efforts in quantum computing is to obtain quantum devices that can be used as the basis for large quantum computers. Many different physical systems are being investigated for their utility in obtaining such devices. Examples include systems based on quantum dots, Josephson junctions, and ion traps. Although we are still a long way from realizing useful quantum computations, the physical system currently closest to the desired goal is based on ion traps.

In order to determine the suitability of a device for quantum computing and to compare it to other devices, it is useful to perform standard “benchmark” experiments. These can be used, for example, to characterize the error in the physical implementation of a quantum algorithm or to demonstrate a process or protocol useful for scaling up quantum computers. Well-known instances include quantum teleportation, quantum error correction, the quantum Fourier transform, entangled state preparations and entanglement purification. Each has been implemented for ion traps in the NIST Physics Lab; here, we focus on the last two, which allowed us to explore procedures involving up to six qubits.

Ion trap quantum computing. To build a quantum computer, it is necessary to supply qubits and the means for (a) preparing them in an initial state, (b) manipulating and coupling them, and (c) measuring their final state. In implementations featuring ion traps, the qubits are represented by the internal states of ions trapped along an axis by means of radiofrequency and static electric fields. For our benchmarking experiments, the ions are Beryllium, ${}^9\text{Be}^+$, cooled to their motional ground state by means of laser cooling, spe-

cifically Raman sideband cooling. The qubit states are given by two hyperfine levels that are manipulated by stimulated Raman transitions. Coupling between qubits is accomplished by use of common vibrational motions. Measurement is by detecting fluorescence: Only one of the two qubit states fluoresces when probed with appropriately tuned and polarized light. A picture of the electrodes used to trap the ions is shown in Fig. 15. The trap features multiple trapping zones and allows for moving, separating and recombining ions. These features are crucial for scaling up the ion trap system for building large quantum computers, and they enabled the implementation of the purification protocol.

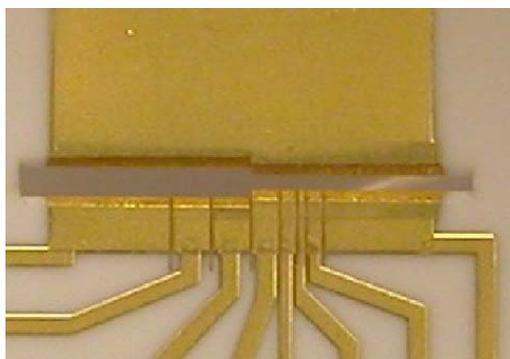


Figure 15. The multizone ion trap used in our experiments. The ions are trapped in the slot midway between the gold electrodes.

Schrödinger cat states. Schrödinger cat states are equal superpositions of a pair of highly distinguishable configurations. They are named for Schrödinger’s (in)famous thought experiment involving a cat in an isolated deadly contraption, which according to quantum theory exists in a superposition of being dead and being alive until “observed”. Because of the distinguishability of the two configurations in a cat state, one expects an equal probability of finding the system in one of the two configurations rather than a superposition. Superpositions are distinguished from probability distributions by the possibility of interference. Because of the fragility of these effects, their demonstration is a good test of the ability to control and couple the qubits.

A single ion qubit’s state space is spanned by configurations or states that can be labeled $|\text{up}\rangle$ and $|\text{down}\rangle$, distinguishable by their spin and energy. However, because in the absence of light they are largely decoupled from the environment, they are not considered highly distinguishable on their own. To create cat states, we use up to six ion qubits. The six qubits have two states $|\text{up}_6\rangle$ and $|\text{down}_6\rangle$, where in state $|\text{up}_6\rangle$ ($|\text{down}_6\rangle$) all qubits are in state $|\text{up}\rangle$ ($|\text{down}\rangle$, respectively). This pair of states is in a sense maximally distinguishable for six qubits and their equal superposition is a cat state of six qubits.

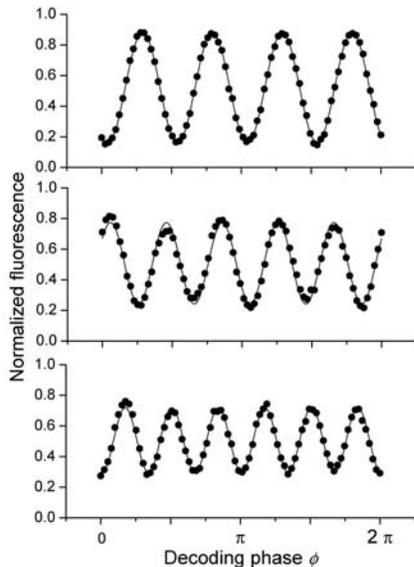


Figure 16. Cat state coherences. The quality of the prepared states is determined by the contrast, i.e., the amplitude of the fringes in an interference experiment. These are shown for cat states of four, five and six ion qubits (top to bottom). The reduction in amplitude is due to the increasing fragility of the states as more qubits are involved.

To determine how well we can create ion qubit cat states, we performed interference experiments that are sensitive only to the coherence in the superpositions. See Fig. 16. The amount of interference (the “contrast”) can be used to determine the fidelity with which we prepared the cat states. The fidelities decreased from 0.76 for four qubits to 0.51 for six, which is state-of-the-art for accuracy of control of this many qubits. An additional property of the cat states is that they have entanglement that cannot be attributed to a proper subset of the qubits, that is, they are globally entangled. The data shows that such entanglement was present.

Purification. One of the most useful resources in quantum computation and communication are pairs of qubits in maximally entangled states. Equivalent to cat states for two qubits, these can be used for teleporting quantum states over long distances without transporting qubits, and for implementing computations involving remote qubits. Because in typical applications the two qubits of an entangled pair are far apart, it is difficult to prepare the state with low error. Entanglement purification is a procedure to reduce the error without further use of quantum communication. It requires the use of local quantum gates and a measurement to compare two pairs of entangled but noisy qubits. If the comparison is “good”, one pair of less noisy qubits is obtained. Entanglement purification will likely play an important role in any large scale quantum computer.

Implementing entanglement purification using the ion trap required four ion qubits. For initial state preparation and the crucial qubit comparison gates, we realized specially designed four qubit gates. To measure the

ions, we separated them into two sets of two, measured the first two at once to determine whether the comparison was good, and then used tomographic state analysis to see how good the entangled state on the remaining two ions was. The purified entangled pair was compared to the unpurified pairs to see whether we were able to improve the entanglement. Because of the complexity of the experiments, errors introduced by the purification procedure prevented an ideal improvement, though it was nevertheless significant (Fig. 17).

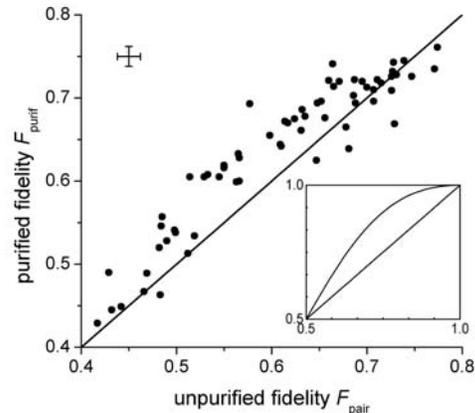


Figure 17. Purified fidelity as a function of unpurified fidelity. The dots in the figure represent individual experiments. Each involved applying a known error to the initially prepared entangled pairs of qubits. The fidelity of the unpurified (purified) pairs determines the x-axis (y-axis, respectively) position. The diagonal line is the line where the purified and unpurified fidelities are the same. There is a range of errors for which an improvement in fidelity is obtained. For comparison, the curved solid line in the inset shows the theoretical purified fidelity as a function of the unpurified fidelity assuming all gates and measurements are ideal.

References

- [1] D. Leibfried, E. Knill, S. Seidelin, J. Britton, R. B. Blakestad, J. Chiaverini, D. B. Hume, W. M. Itano, J. D. Jost, C. Langer, R. Ozeri, R. Reichle and D. J. Wineland, Creation of a Six-Atom “Schrödinger Cat” State, *Nature* **438** (2005), pp. 639-642.
- [2] R. Reichle, D. Leibfried, E. Knill, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, Experimental Purification of Two-Atom Entanglement, *Nature* **443** (2006), pp. 838-842.

Participants

E. H. Knill (MCS D); D. Leibfried, D. Wineland, and colleagues in the Ion Storage Group (PL)

OOF: Object-Oriented Finite Elements

Rapid increases in computer power have created an opportunity for investigators to use computational methods to tackle both larger and more complex problems than had been possible previously. However, this will occur only if we can create tools which take advantage of this computational power, while at the same time encapsulating the sophisticated numerical methods which handle the complexity. In the field of materials science, the direct computational study of structure-property relationships in complex microstructures is enabled by OOF, the Object-Oriented Finite Element system. A collaboration between MCSD and the NIST Materials Science and Engineering Laboratory, this new tool provides finite-element computational capabilities in a modular, extensible framework designed to speak the language of materials scientists. Version 2.0 of OOF was released to the public in September 2006.

Stephen Langer

The Object-Oriented Finite Element (OOF) project produces software that enables the direct computational study of structure-property relationships in materials with complex microstructures. When running OOF [1], a user assigns material properties to the features in an image of a material's microstructure, and then performs virtual experiments on the material. Unlike commonly available commercial finite element codes, OOF features both a powerful suite of tools for adapting a finite-element mesh to the microstructural geometry of an image, and a modular and extensible scheme for adding new or customized property data to the underlying model. Property data can also be specified independently of a particular problem, and shared between problems and between researchers.

The archetypal OOF user is a materials scientist with deep knowledge of the materials under study, but possibly little or no knowledge of the more technical aspects of computational modeling. Such an investigator is interested in examining multiple interactions between regions of a spatially complex microstructure, but is not an expert in modeling techniques, finite element basis functions, sparse matrix solvers, etc. The OOF program provides, on the one hand, a powerful tool set for managing the parts of the problem where the investigator is knowledgeable (e.g., constitutive rules, microstructural geometry) and, on the other hand, push-button access to sophisticated numerical tools which implement the model. As software, OOF occupies a niche unlikely to be served by commercial software, and benefits substantially from the breadth of

expertise and institutional memory present at NIST. The development team combines expertise in physics, materials science, and software engineering.

Inputs to an OOF analysis are one or more images representing the microstructure, and the physical properties of the components of the microstructure. The images can be micrographs obtained experimentally or the output of a simulation. The user segments each image into homogeneous regions, assigning to each appropriate material properties. Using the adaptation tools, a finite element mesh is constructed which matches the geometry of the microstructure. Boundary conditions are applied, representing loads in a virtual experiment, the equations are solved, and the resulting configuration of fields and fluxes is measured.

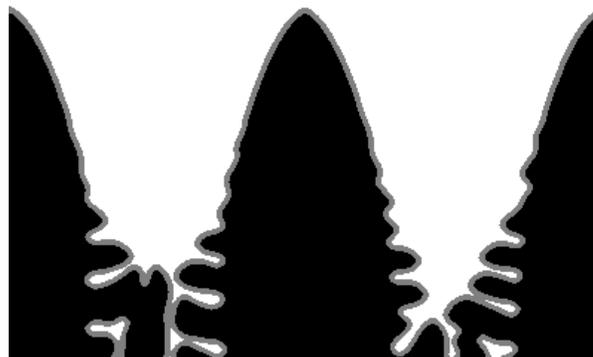


Figure 18. A microstructure from a simulation of dendritic growth. The three colors correspond to regions with different chemical compositions. (Courtesy of J. Warren, NIST MSEL.)

This very general scheme not only allows for a direct simulation where real-world properties and microstructure are known, but also allows OOF to function as a design tool. If properties are known, it's possible to use OOF to "search" among candidate microstructures for optimal behavior. If, on the other hand, the microstructure is known but the properties are not, a parametric study can be done to determine optimal parameters, either to reproduce observed behavior or to achieve desired behavior. The relative rapidity and cheapness of the OOF analysis enables many more design iterations than could be achieved by direct experimentation.

Earlier versions of the OOF software have already shown the value of this approach. Version 1.0, originally released in 1998, has been modified and extended many times since. In 1999, it earned a "Technology of the Year" award from *Industry Week* magazine. However, OOF1 failed to realize one of the major goals of the project, i.e., to create a software platform that users could easily extend by adding their own physical models. The original version of OOF1 solved only linear

elasticity problems. A second version was created at NIST to solve problems with coupled elasticity and thermal diffusion, and an MIT student created versions that incorporated piezoelectricity and chemical diffusion. These were used to study ceramic thermal barrier coatings on turbine blades [2], Li-ion batteries [3], and polycrystalline piezoelectric transducers [4], among other things [5, 6]. In the process it became clear that OOF1 was too difficult to extend, and that maintaining a zoo of different versions was not practical.

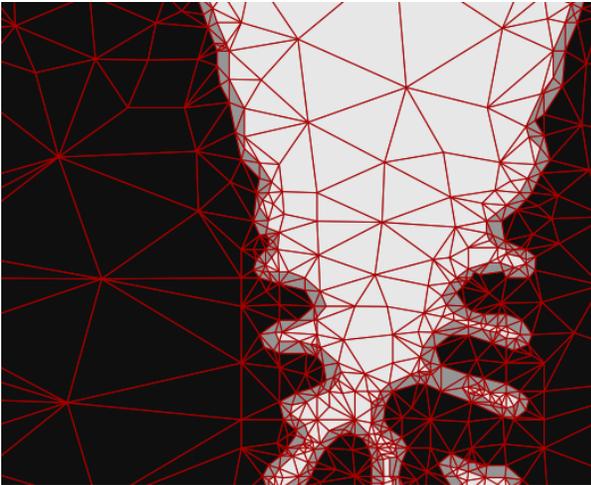


Figure 19. A portion of the microstructure, meshed with OOF2.

OOF2 was initially released in early September 2006. It is a complete rewrite of the program designed to overcome the shortcomings of OOF1. The features of all the different versions of OOF1 coexist amicably within the OOF2 framework. Unlike OOF1's fixed material types, OOF2's materials are created by assembling a collection of diverse properties (e.g., elastic modulus, piezoelectric modulus, heat capacity, orientation) at run time. New property definitions can be created and loaded in a modular fashion. In addition to more powerful and flexible ways of defining material properties, OOF2 allows more kinds of finite elements. While OOF1 only supported linear triangles, OOF2 contains linear and quadratic triangles and quadrilaterals, and new types can be added easily. OOF2 also contains improvements in mesh creation, boundary conditions, and solvers. OOF2 is threaded, meaning that it can perform more than one operation at a time, which vastly improves the user experience.

At the OOF workshop held at NIST in August 2006, some 40 OOF users shared their experiences with OOF1, and were introduced to OOF2. Users came from five US industries, three US government laboratories, and 18 universities, representing a broad cross-section of instructional, research, and commercial interests.

Current and future developments will enhance OOF2 in a number of ways. The development team is currently

adapting the software to run in parallel on a cluster, and is also extending its range to include periodic boundary conditions, plasticity and other nonlinear properties, time dependence, and partitioning of problems into multiple subproblems. Parallelization will enable the software to handle large data sets, and is a prerequisite for a planned extension to a fully three dimensional computational domain. Further expansion beyond these broad capabilities will be driven by user requirements.

References

- [1] S. A. Langer, E. R. Fuller, Jr., and W. C. Carter, OOF: An Image-Based Finite-Element Analysis of Material Microstructures, *Computing in Science and Engineering* **3**:3, 15-23 (2001).
- [2] C. H. Hsueh, P. F. Becher, E. R. Fuller Jr., S. A. Langer, and W. C. Carter, Surface-Roughness Induced Residual Stresses in Thermal Barrier Coatings: Computer Simulations, *Materials Science Forum* **308**, 442-449 (1999).
- [3] R. Edwin Garcia, Yet-Ming Chiang, W. Craig Carter, Pimpa Limthongkul, Catherine M. Bishop. Microstructural Modeling and Design of Rechargeable Lithium-Ion Batteries, *Journal of the Electrochemical Society* **152**, A255-A263 (2005).
- [4] R. Edwin Garcia, W. Craig Carter, and Stephen A. Langer, The Effect of Texture and Microstructure on the Macroscopic Properties of Polycrystalline Piezoelectrics: Application to BaTiO₃ and PZN-PT, *J. Am. Ceramics Society* **88**:3 750-775 (2005).
- [5] H. Singh, A. M. Gokhale, Y. Mao and J. E. Spowart, Computer Simulations of Realistic Microstructures of Discontinuously Reinforced Aluminum Alloy (DRA) Composites, *Acta Materialia* **54**, 2131-2143 (2006).
- [6] Thomas Weiss, Siegfried Siegesmund, and Edwin R. Fuller, Jr., Thermal Degradation of Marble: Indications from Finite-Element Modelling, *Building and Environment* **38**, 1251-1260 (2003).

Participants

S. Langer (MCS), A. Reid (Drexel), Seung-Il Han (UMBC), R. Lua (Penn State), V. Coffman (Cornell), E. Garcia (Purdue)

<http://www.ctcms.nist.gov/oof/>

Modeling Fluid Flow in Materials Processing

Materials processing techniques often involve the initial preparation of a sample of material in one thermodynamic state, such as the material's liquid phase at a high temperature, and carefully allowing the sample to transform into a desired final state, say the sample's solid phase at lower temperatures. During this process an interface separates the material that remains in the initial state from the material that has reached its final state. The shape of this interface has a critical effect on the quality of the resulting material. In many applications a planar interface is optimal, which leads to the formation of homogeneous samples. Such samples will typically have desirable mechanical or electronic properties. In contrast, nonplanar interfaces can lead to unfavorable inhomogeneities that can degrade the performance of the material. It is often found that while a planar interface is possible in theory, in practice it may become unstable under the combined effects of convection and heat transfer. Determining the stability of the interface is therefore critical in order to understand and control materials processing in many applications.

Here we consider the processing a sample from one liquid phase to another, such as in the condensation of a liquid from its gas phase. In this case we need to consider the stability of a fluid-fluid interface subject to a temperature gradient perpendicular to the interface. For a phase-transforming system this results in the generalization of a classical hydrodynamic stability problem for fluids that are subject to the effects of buoyancy and surface tension gradients. In such systems, for large enough temperature gradients, a planar fluid-fluid interface can undergo bifurcations to nonplanar geometries with associated changes in the state of convection and heat flow throughout the sample.

Geoffrey B. McFadden

The study of the stability of a fluid-fluid interface is important in a number of scientific and technological applications. In this project we consider two fluid layers separated by a horizontal planar interface subject to a vertical temperature gradient. The effects of various driving forces on the stability of the system can be taken into account, including buoyancy (known as Rayleigh-Benard convection), the effects of bulk density differences (known as Rayleigh-Taylor instabilities), and the effects of surface tension gradients along the interface (known as Marangoni instabilities). If the two layers represent different phases of the same material the stability results for a two-phase bilayer system are quantitatively and even qualitatively different than

for those of an immiscible system. For the two-phase system we find particularly interesting stability results at low wavenumbers, where we discover a new mode that is sensitive to the differences in material properties, specifically the enthalpy and entropy difference between the two phases.

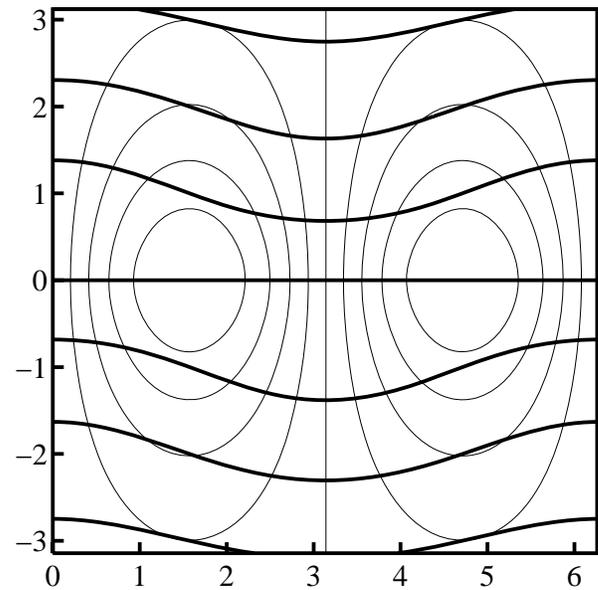


Figure 20. Streamfunction contours (light lines) and temperature contours (dark lines) for flow in the two-layer system. The interface corresponds to the streamline at $y = 0$.

To estimate the relative importance of these types of instabilities for a two-phase system, we consider a bilayer geometry in which a horizontal fluid-fluid interface separates two semi-infinite layers of a single-component fluid. This problem is still sufficiently complex that we generally resort to a numerical determination of the linear stability of the system, including the possibility of temporally oscillatory modes with significant interfacial deformations. The governing equations contain a large number of dimensionless parameters, including a Rayleigh number, Marangoni number, Bond number (ratio of the forces of buoyancy and surface tension) and Crispation number (a measure of interface deformation), and ratios of material properties of the two phases.

In particular, we have performed linear stability calculations for horizontal fluid bilayers that can undergo a phase transformation, taking into account both buoyancy effects and thermocapillary effects. We compare the familiar case of the stability of two immiscible fluids in a bilayer geometry (the “inert” case) with the less-studied case in which the two fluids represent different phases of a single-component material. The two cases differ in their respective interfacial boundary

conditions. We find that, as in the inert case, the two-phase system can be linearly unstable to either heating from above or below. Instability in the two-phase case persists to small wavenumbers in the case of heating from below.

To help understand the mechanisms driving the instability on heating from below, we performed both long-wavelength and short-wavelength analyses of the two-phase system. The short-wavelength analysis shows that the instability is driven by a coupling between the flow normal to the interface and the latent heat release at the interface. The coupling is illustrated in the contours of the temperature and streamfunction for the two-phase system case; see Fig. 20. Here the fluid flow is normal to the interface, and the downflow in the center tends to compress the distance between isotherms near the interface in the upper phase, while expanding the distance between them in the lower phase. The opposite is true for the regions with upflow at the interface. The resulting net change in the temperature fluxes at the interface is balanced by the evolution of latent heat at the interface, which in turn is driven by the vertical velocity at the interface. When the system is heated from below as shown, these effects reinforce each other to drive the instability. When the system is heated from above, the effects are in opposing directions and no instability is possible.

The mechanism for the large wavelength instability is more complicated; see Fig. 21. The detailed nature of the instability depends on the Crispation and Bond numbers. We find that a two-phase system that is heated from below is subject to a type of morphological instability that is similar in origin to other, more familiar, instabilities that occur in materials processing applications. Multiphase systems are typically subject to instability if one or more of the phases is in a thermodynamic state of metastability; specifically, if there are regions of space where the actual phase of the system is not the phase of lowest free energy under the given conditions. Examples include *superheating* (a solid above its melting point), *supercooling* (a liquid below its melting point), *supersaturating* (a solution beyond its solubility limit), or compressing a gaseous component beyond its vapor pressure. A crystal of binary alloy growing from its melt is subject to instability if the liquid ahead of the interface is *constitutionally supercooled*. In this case, the driving force for instability is the release of free energy accompanying the phase transformation from a higher energy phase to a lower energy phase. The resulting instability generally displays a wavelength that is determined by a balance between surface energy (stabilizing) and the difference in bulk free energy between the phases (destabilizing).

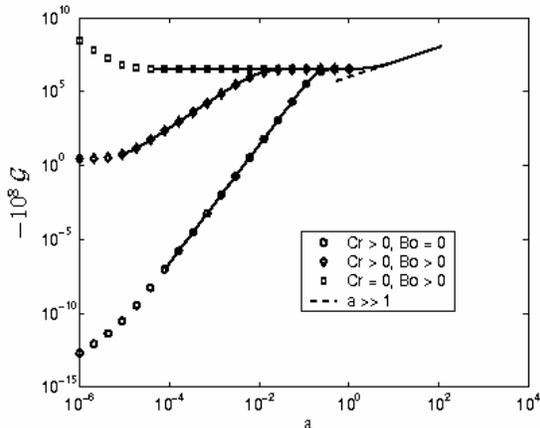


Figure 21. Marginal stability curves for a two-phase water-steam system heated from below. The vertical axis represents critical values of the dimensionless temperature difference across the layers, and the horizontal axis is the lateral wavenumber of the instability.

Three cases are shown in Fig. 21. The upper curve results from setting the Crispation number to zero while the lower curve corresponds to setting the Bond number to zero. Both the Crispation number and Bond number are present for the middle curve. The solid curves correspond to numerical results. The symbols on the curves correspond to analytical results from a small-wavenumber approximation. The results illustrate the sensitivity of the system to the effects of the Bond number and Crispation number.

This work is currently being extended to the important case of multicomponent materials.

References

- [1] G. B. McFadden, S. R. Coriell, K. F. Gurski and D. L. Cotrell, Onset of Convection in Two Liquid Layers with Phase Change, *Journal of Fluid Mechanics*, submitted.

Participants

G. B. McFadden (MCSD); S. Coriell (MSEL); K. Gurski (George Washington University); D. Cotrell (Lawrence Livermore National Laboratory)

An Accurate Model for Measuring Fluorophore Photodegradation Dynamics

Fluorescent dyes, stains and probes play a prominent role in the measurement and detection of scientific phenomena on a small scale. Difficult problems in material science, chemistry, and nanotechnology can be investigated, for example, by using fluorescent techniques to enable the observation of the dynamics of single molecules. In biomedical and biotechnology applications fluorescent materials are used to obtain qualitative and quantitative information about processes on the cellular and sub-cellular level. However, the sensitivity and accuracy of such techniques is severely limited by natural fluorescent signal decay. This process, called photodegradation, is the sequence of photochemical reactions that transform excited fluorescent particles (fluorophores) to a non-fluorescent species. A thorough understanding of fluorophore photodegradation dynamics is necessary to enable the effective use of this measurement technique. In this work, we develop a mathematical model of photodegradation dynamics that justifies the use of a particular functional form to fit fluorescent signal measurements in the frequency domain.

Fern Y. Hunt

Photodegradation is due to a combination of processes that occur on multiple time scales, which makes the experimental interpretation of the time-dependent dynamics difficult. Recently scientists in the Biochemical Sciences Division of the NIST Chemical Sciences and Technology Laboratory (CSTL) developed a frequency domain based method for measuring photodegradation using an apparatus known as a lock-in amplifier [1]; see in Fig. 22. A sample of fluid containing fluorophores flows vertically in the x direction through a flow cell and is stimulated by laser light shining through the center. The fluid motion allows fresh populations of unperturbed molecules to be exposed to the light and minimizes the effects of uncontrollable convective and diffusive mass transport. When the intensity of the light fluctuates periodically the fluorescent emission is periodic with the same frequency, but it is shifted in phase. The emission is converted to an electrical signal, and by suitable signal multiplication and filtering one can reduce noise and extract a weak signal. The experiment is repeated for a number of excitation frequencies, thus allowing the estimation of the fluorescent photodegradation rate from the parameters of an impedance curve that is fitted to the measurements.

Specifically, the link between the measured phase shift and the photodegradation rate is a formula for the total fluorescence in the apparatus as a function of time. The formula comes from a mathematical model of the underlying process, allowing the fluorescence to be written in terms of physical parameters [2]. Here, the light is assumed to be distributed uniformly throughout the flow cell. In this case, the total fluorescence at any time can be written in terms of the averaged fluorophore concentration. If the impedance curve (showing the ratio of the imaginary and real parts of the Fourier transform of the fluorescence as a function of laser light frequency ω) is fit to the functional form

$$\frac{p_1 \omega}{p_2 + \omega^2}, \quad (1)$$

the parameters p_1 , p_2 can then be expressed in terms of the photodegradation rate and other physical parameters thanks to the tractability of the expression for the total fluorescence.

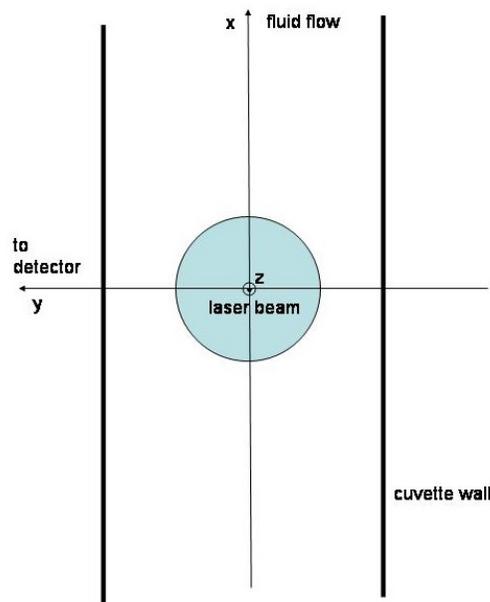


Figure 22. Schematic diagram of a lock-in amplifier for measuring fluorophore photodegradation.

In collaboration with CSTL scientists, we have examined the more physically realistic case of a laser beam that has a Gaussian power distribution. Here the photochemical reaction rates, which depend on power, will also depend on the location in the beam. Therefore the time evolution of the total fluorescence has to take into account the fluorescent response in a non-uniformly illuminated region.

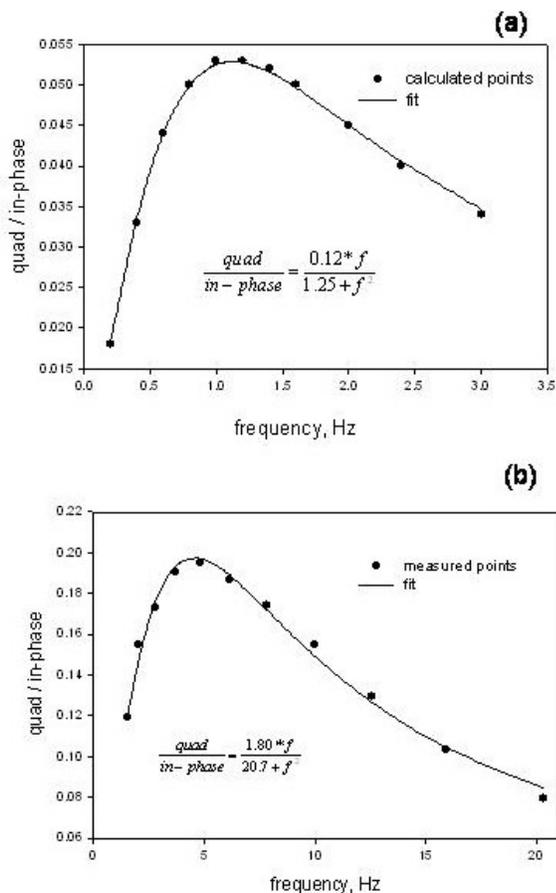


Figure 23. Impedance data for (a) the model and (b) physical measurement. Data points are fit to the same functional form in each case, yielding the shown curves.

Starting with a (rather simplified) model of the kinetics associated with excitation, relaxation and photodegradation, we analyzed a pair of coupled partial differential equations that describe the changes in fluorophore populations due to the kinetics and the convective fluid flow. The boundary conditions for this problem take into account the continuous entry of fluid into the apparatus from its base. Using singular perturbation techniques, we obtain a reduction in the mathematical complexity of the model by taking advantage of disparate time scales (photodegradation is much slower than any other process in the model). We then derive a single partial differential equation for the total fluorophore population (to leading order in a perturbation parameter). When the amplitude of the periodic component of the light is relatively small, spatial averaging can be performed to yield an ordinary differential equation that is accurate up to a small error. The measured fluorescence can be modeled approximately in terms of the spatially averaged fluorophore population. These two equations, then, constitute an experimentally accessible model for the fluorescence that can be used to estimate the photodegradation rate in a manner analogous to [2].

Fig. 23(a) shows the impedance curve generated by the solution of the model equations, while Fig. 23(b) shows the curve obtained by measurements with the lock-in amplifier. While they demonstrate only qualitative agreement, both sets of data are fit well with the functional form of Eq. 1.

Our analysis is comprehensive. The steps used to derive the model are mathematically justified and the order of magnitude of the errors in the perturbation approximation is estimated. In addition to being applicable to either the uniform, Gaussian, or Gaussian-like light distribution, we believe our analysis would apply to the more complicated kinetic model that provides a more detailed description of the photodegradation process. Two technical papers describing this work are in process [3, 4]. Using analytical techniques in [4], we can obtain information about the relationships between the fit parameters and the physical constants of the experiment. This will allow us to obtain quantitative agreement in the impedance curves.

References

- [1] A. K. Gaigalas, L. Wang, and R. F. Vogt Jr., Frequency-domain Measurement of the Photodegradation Process of Fluorescein, *Photochemistry and Photobiology* **76** (2001), pp. 22-28.
- [2] A. K. Gaigalas, L. Wang, K. D. Cole, and E. Humphries, Photodegradation of Fluorescein in Solutions Containing n-Propyl Gallate, *Journal of Physical Chemistry A* **108** (2004), pp. 4378-4384.
- [3] F. Y. Hunt, A. K. Gaigalas, and L. Wang, Mathematical Foundation of the Frequency Domain Technique for Measuring Photodegradation, *Journal of Physical Chemistry A*, submitted, 2006.
- [4] A.K. Gaigalas, L. Wang and F.Y. Hunt, Harmonic Solution to the Kinetic Equations Describing Photoreactions in a Gaussian Laser Beam, in preparation, 2006.

Participants

F. Y. Hunt (MCS D); A. K. Gaigalas, L. Wang (CSTL)

Numerical Optimization of Complex Instrumentation

A novel approach has been created for the selection of optimal instrument parameters that yield a mass spectrum which best replicates the molecular mass distribution of a synthetic polymer. The application of implicit filtering algorithms was shown to be a viable method to find the best instrument settings while simultaneously minimizing the total number of experiments that need to be performed. This includes considerations of when to halt the iterative optimization process at a point when statistically significant gains can no longer be expected. This work represents part of an effort to develop an absolute molecular mass distribution polymer Standard Reference Material by matrix-assisted laser desorption/ionization time-of-flight (MALDI-TOF) mass spectrometry.

Anthony J. Kearsley
William E. Wallace

Typical analytical instrumentation optimization is performed by the analyst by simply applying the “factory settings” or by “optimizing by eye”. This is because an exhaustive search of the parameter space for modern instrumentation with many adjustable parameters is prohibitively time consuming. Nevertheless, a variety of mathematical methods exist that can allow the experimentalist to optimize instrument settings without performing an exhaustive search. Broadly classified, these methods are all forms of numerical optimization. When the topology of the search space is very complex, for example when it has great sensitivity to one or more parameters (as mass spectrometers often do) the methods used are part of the field of nonlinear programming. They are nonlinear because some (or all) of the instrument parameters do not have a linear relationship between parameter value and measurement response. An example is laser intensity in matrix-assisted laser desorption/ionization time-of-flight (MALDI-TOF) mass spectrometry and its effect on signal to noise ratio, where a relatively sharp threshold is observed experimentally.

Stochastic numerical optimization methods are important in mass spectrometry because all mass spectra have random noise. This noise varies as the instrument parameters are adjusted, and the noise will often change across the spectrum. Measurement noise presents a significant challenge to any optimization method, especially for cases where signal to noise is not the measurand to be optimized. Nevertheless, numerical optimization methods offer experimentalists a way to tune the instrument parameters to achieve the desired goal without having to search all possible pa-

rameter combinations. To measure the absolute molecular mass distribution of a synthetic polymer, it would be ideal to locate a region in parameter space where the instrument response function was uniform across the entire mass range. Finding the instrument response function is necessary to calibrate the intensity axis of the mass spectrum, that is, to go from mass spectrum to molecular mass distribution. If the instrument response function is uniform then the relative peak areas in the mass spectrum correspond directly to the relative abundances of individual n mers in the sample. A uniform instrument response function would be a line of zero slope. If not uniform, the instrument response function could slowly vary across the mass range, preferably linearly with mass. The optimal conditions are those that give the simplest (or flattest) instrument response function, that is, the one with the smallest derivative.

To measure the instrument response function, a gravimetric mixture was made of three low polydispersity polystyrenes that were very close in average molecular mass. The optimal instrument settings are those that provided the closest match between the total integrated peak intensity of each of the three polymers with the known gravimetric ratios. Note that there is no guarantee (or even assumption) that the optimal instrument settings that give the flattest instrument response function will also yield optimal signal to noise ratios. In fact, there is no reason to believe that a search for the instrument settings that optimize the response function will not lead into a region where the mass spectra becomes so noisy as to make peak integration impossible. Thus, to find the optimal instrument settings we used stochastic gradient approximation methods. These methods have proven to be extremely robust in cases where the measured data is very noisy.

Optimization is performed by defining an objective function $J(x)$ where x is a vector consisting of the instrument parameters. In our case, the objective function is the sum of the squared differences between the amount of each polymer in a mixture created gravimetrically and the amount of each polymer in the mixture found by mass spectrometry. When this function is zero, the gravimeter concentrations match the concentrations found by mass spectrometry and the instrument is optimized. The function $J(x)$ is a noisy function with respect to the parameter vector x due to the inherent statistical noise in the mass spectra. This complicates the task of numerically locating the minimum of $J(x)$. The fact that each evaluation of $J(x)$ requires an experiment, and subsequent interpretation of experimental results, means that there is a high cost for each function evaluation. This further complicates any nu-

merical procedure that seeks to minimize $J(x)$. Finally, there are values of the vector x (for example, out of range instrument parameter settings) for which $J(x)$ cannot be evaluated.

One method for minimizing noisy functions that seeks to approximate the gradient of the objective function is called implicit filtering. Broadly speaking this method uses a very coarse grained step-length to build a finite difference approximation to the gradient of $J(x)$. This gradient is then used to generate descent directions for a minimization process. As iterates draw closer to the solution, and the objective function decreases, the finite difference step-length is decreased until it approaches a number small enough to suggest convergence to a minimum.

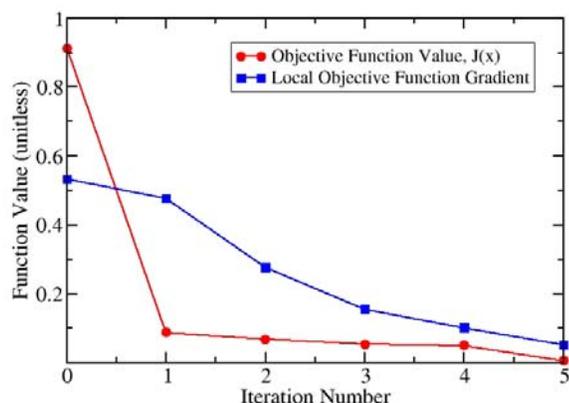


Figure 24. The objective function $J(x)$ and its local gradient value as a function of iteration step

Fig. 24 shows the objective function $J(x)$ and its local gradient value as a function of iteration step. There is an initial steep drop in the objective function followed by gradual movement to the optimal parameter settings. The gradient of the objective function also decreases steadily. These monotonic responses indicate that the optimization routine is stable. At the optimum value the objective function is so small that it cannot be reduced further due to the inherent noise in the measurement. Likewise, the step size indicated for each parameter at this point is so small as to be below the precision of the instrument's settings.

Fig. 25 shows the individual instrument parameter values as a function of iteration number. The values oscillate about their final values as the optimization proceeds. The laser intensity undergoes the greatest excursions decreasing in the first two iterations, returning to its initial value in the third iteration, and then increasing in the fourth iteration before settling into its final value. The four other parameters make an excursion in the direction of their final values in the first iteration, return to their initial value in the second iteration, and find the equilibrium values by the third iteration. This zigzag pattern is characteristic of the

nonlinearity of the system. The nonlinearity arises from the fact that the instrument parameters are coupled, that is, varying one requires all others to vary in response if $J(x)$ is to move closer to its optimal value. Thus, the vector x_{i+1} has a tendency to be normal to the vector x_i . The laser intensity varies the most and seems to be the dominant variable. It seeks its stable value before the other parameters can settle down.

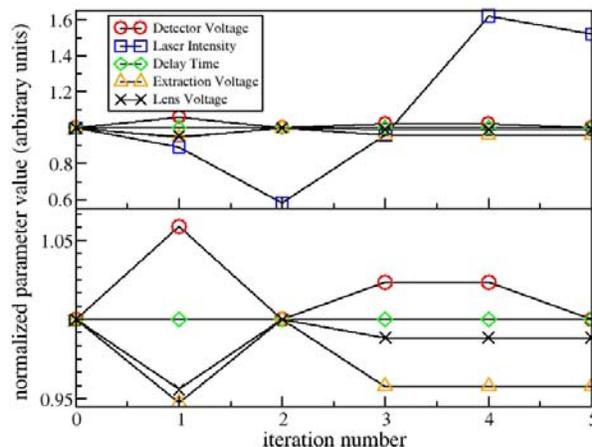


Figure 25. Individual instrument parameter values as a function of iteration number

Summary. A specialized noise-adapted filtering method has been applied to the problem of finding the optimal instrument parameters for a MALDI-TOF mass spectrometer. Finding the optimal instrument parameters was a critical step in creating an absolute molecular mass distribution polymer Standard Reference Material. The task of tuning the instrument's five main parameters could not be approached by exhaustive search methods given the amount of effort needed to take and to reduce the data in a statistically meaningful way at each set of instrument parameters. Additionally, this method produces an estimate of the sensitivity of each optimal parameter estimate not available to traditional exhaustive search methods. Each of the subtasks in the process could be automated to create an integrated closed-loop optimization scheme.

References

- [1] W. E. Wallace, K.M. Flynn, C.M. Guttman and A. J. Kearsley, Numerical Optimization of Matrix-Assisted Laser Desorption/Ionization Time-of-Flight Mass Spectrometry: Application to Synthetic Polymer Molecular Mass Distribution Measurement, submitted.

Participants

A. J. Kearsley (MCSD), K. M. Flynn, C. M. Guttman, W. E. Wallace (MSEL).